RISK ASSESSMENT

For

Central Web Tomcat Service

Minor Application

[ID ####]

Responsibility for the Minor Application and its operation as described in this plan is accepted by:

_____ Date:_____

Service Owner - Peter J. Rzeminski II

_____ Date:_____

ESO Department Head – Michael Rosier

_____ Date:_____

Chief Information Security Officer – Irwin Gaines

_____ Date:_____

CCD Division Head – Jon Bakken

Contents

Information and Security Contacts:

Title	Name	email	Telephone	Initials
Service Owner (s)	Peter J. Rzeminski II	ptr@fnal.gov	630.840.5524	
System Managers – Apache httpd	Andrew Duranceau John Inkmann	adurance@fnal.gov inkmann@fnal.gov	630.840.6457 630.840.6508	
System Managers – NAS	Andrew Romero	romero@fnal.gov	630.840.4733	
System Managers – Linux OS	James O'Leary	joearly@fnal.gov	630.840.2230	
System Managers – Virtual Environment	Briant Lawson	blawson@fnal.gov	630.840.2944	
Management Contact	Jon Bakken (Division) Michael Rosier (Dept) Peter J. Rzeminski II (Group)	bakken@fnal.gov mrosier@fnal.gov ptr@fnal.gov	630.840.4790 630.840.8385 630.840.5524	

SYSTEM IDENTIFICATION

System Name/Title

Fermilab identifier CSP- GSS-#### has been assigned to the system discussed throughout this Risk Assessment and will be referred to as the Central Web Tomcat Service.

System Type

This system is the Central Web Tomcat Service Minor Application (MA) and is contained in the General Computing Enclave.

Responsible Organization

Fermi National Accelerator Laboratory
PO Box 500
Batavia, IL 60510

System Operational Status

It is in the Operational phase of its life-cycle.

General Description/Purpose

The Central Web Tomcat Service is a centrally managed Apache Tomcat Java application server platform capable of serving website content to the Internet using both IPv4 and IPv6 addresses. Its purpose is to provide a stable Java application server platform for all employees, groups, offices, departments, experiments and any other approved entity connected to Fermilab requiring a Java application server.

Introduction

The layout of the web infrastructure is designed to provide maximum stability and uptime for minimal cost. The service uses a pair of virtual servers running Red Hat Enterprise Linux (RHEL) behind the F5 Big Iron load balancer. Content is stored on the BlueArc and is mounted via NFSv4 to each web server.

The main components of the web service is the F5 Big Iron (F5), the web servers, and the BlueArc.

The flow of web traffic is broken up into a number of layers with each layer designed to provide as much redundancy as possible.

The initial layer is the F5 Big Iron (F5) where web traffic flows in the F5 and is directed to the second layer; the web servers.

The web servers are three components, the active server, the stand-by server, and the Site-Down Service servers. If the active server is unavailable, traffic is directed to the stand-by server. If both the active and the stand-by servers are unavailable, then traffic is directed to the Site-Down Service.

When traffic reaches the web servers, the Apache httpd process identifies Java application requests and proxies to Tomcat. For all other requests, httpd locates the appropriate content stored on the BlueArc, mounted via NFSv4, and presents to the content to the requestor.

Security is maintained by adhering to the given baselines to each layer and by the conservative application of ACL's to the files and directories on the content file system. Access is managed by utilizing existing authentication infrastructure; specifically the UNIX Kerberos real and the Windows FERMI AD realm. Access is granted/removed through tickets submitted through the Service Desk with the approval of the website managers.

This is the Tier 7 category of the Central Web Service. It is similar to Tier 1, except for the filesystem ACLs which are similar to Tier 2 and the ability to execute Java applications. Users will NOT be allowed to SSH into the webserver.

Authentication

There are a number of authentication methods being used with the Central Web Service.

Using a SMB, CIFS, or Windows **Share** connection to the BlueArc, the user must use their FERMI Domain credentials.

To ssh into FNALU, the user must first obtain a valid Fermilab UNIX Kerberos principle.

When accessing content through a web browser, the content owner will have the option of using SAML/SSO via PingFederate to let users authenticate to their website before being able to see content. PingFederate uses SERVICES Domain credentials.

System Description and Boundaries

The central web service infrastructure relies on the following third-party services

- Big Iron F5
    - Managed by the Networking Department, it is used to load-balance inbound traffic
- Virtual Server Infrastructure
    - Managed by the ESO Department / VMS Group, it is used to run the Apache httpd servers
- Red Hat Enterprise Linux (RHEL)
    - Managed by the ESO Department / USS Group, this is the operating system used for the service.
- Apache httpd Software

    - Managed by the ESO Department / WSA Group, this is the software used to present web content to the Internet.
- Apache Tomcat software
    - Managed by the ESO Department / WSA Group, this is the software used to execute Java applications.
- BlueArc File Server (NAS)
    - Managed by the ESO Department, this is the NFS File mount file system where user content and configuration files are stored. It also serves as the primary access point where customers & users will access their content.
- User Content
    - Managed by the customer & user, all content stored within the website's directory is the responsibility of the site managers for the given website.

Virtual Server Design

System Name	Service Tier	Function	Virtual Server Location
Web700#	Tier VII	Standard Service offering for the customers needing to serve Java applications. See the Web Services SLA for the specifics.	FCC2

Webserver Cluster Design (Real World Example)

Cluster Name	Web Servers	Load Balancer Priority	Comments
VIP-WEBT7C01	web7001 131.225.70.27	150	Primary Content Server
	web7002 131.225.70.30	150	Secondary Content Server
	web-sorry01 131.225.70.234	50	Site Down Service – Server #1
	web-sorry02 131.225.70.14	1	Site Down Service – Server #2

Server Tier Offerings (non-exhaustive)

Offering	TIER VII
HTML	YES
JAVA	YES
PERL	YES
PHP	YES
PYTHON	NO
SAML/SSO Authentication to SERVICES via PingFederate	YES
MySQL/Postgres Client Libraries for Perl/PHP/Python etc.	YES
Direct SSH to Content	NO
Any programming language within the YUM repository	Upon Request with Justification
Any software package within the YUM repository	Upon Request with Justification (higher level of scrutiny)

User Access Points

Method	Availability	Port(s)	Source	Access Requirements
HTTP	World Wide	80	Apache httpd web server	None
HTTPS	World Wide	443	Apache httpd web server	None
SAML/SSO	World Wide	80/443	PingFederate	Optional secure authentication method through the browser
SSH **SCP** **SFTP**	Fermilab Network Only	22	Apache httpd web server	Available to Web Systems Administration Group only.
SMB, CIFS, Windows Share	None	139 & 445	BlueArc NAS	Not Available

Internet Traffic Flow

Central Web Service
Internet Traffic Flow

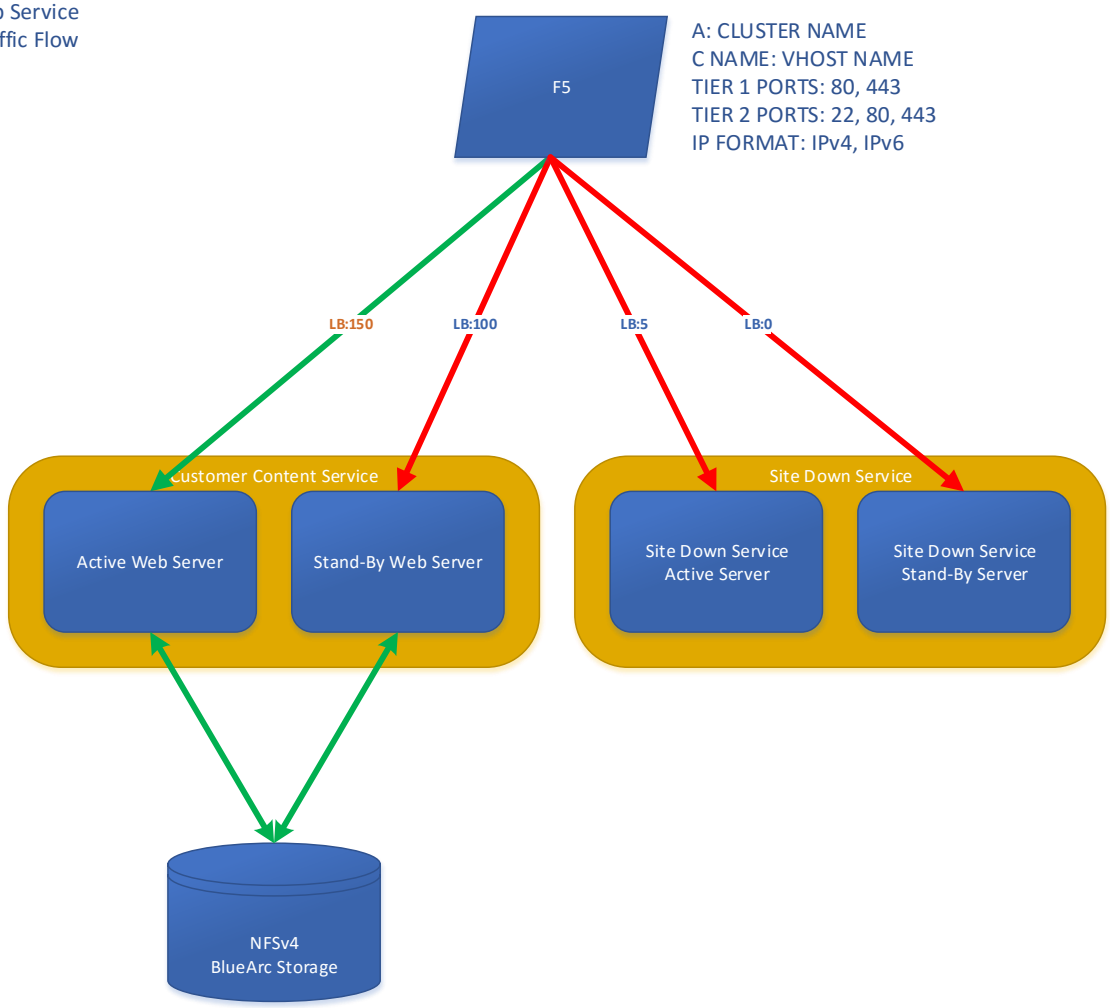

Internet Traffic Flow: User Content Access Methods

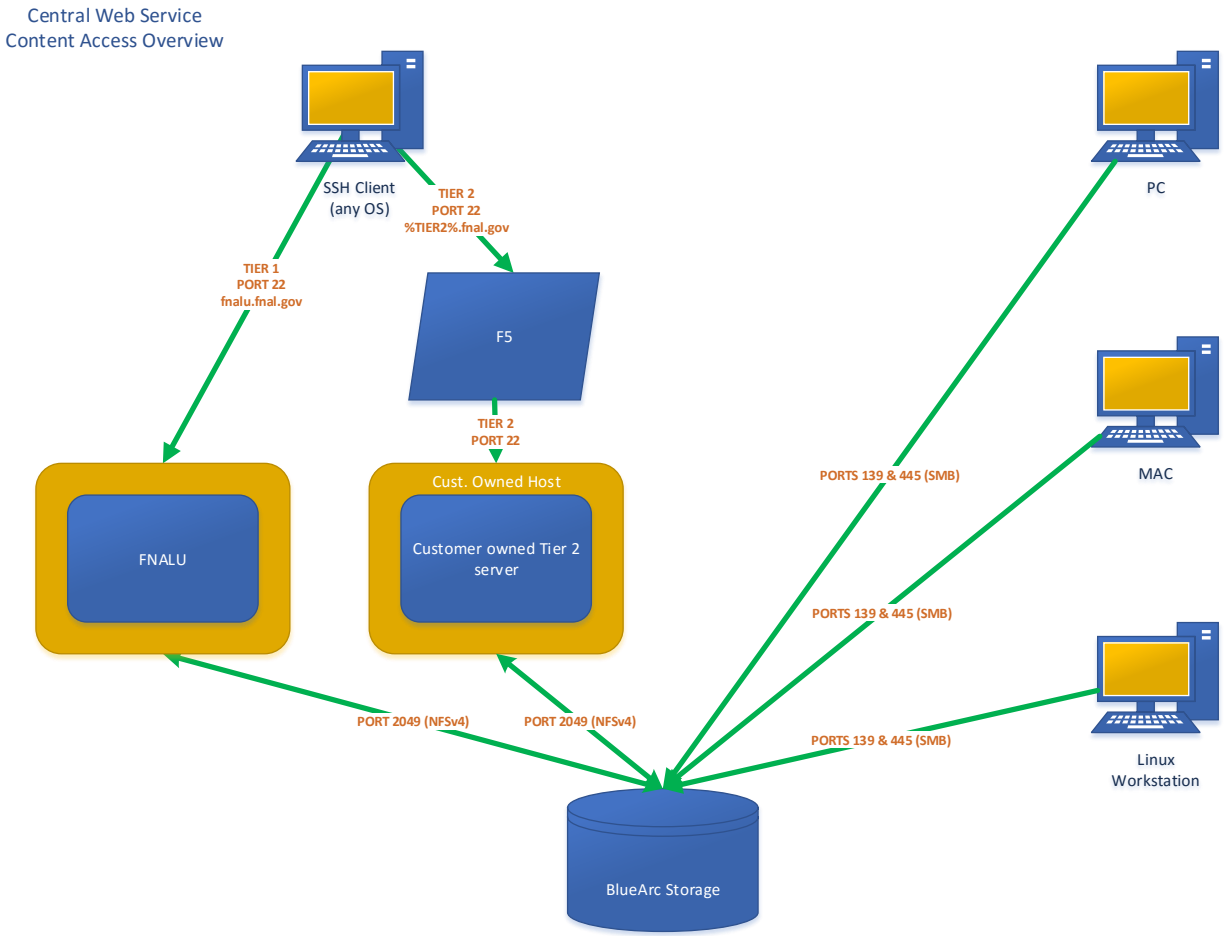

Server Startup: Configuration File Processing & Load Order

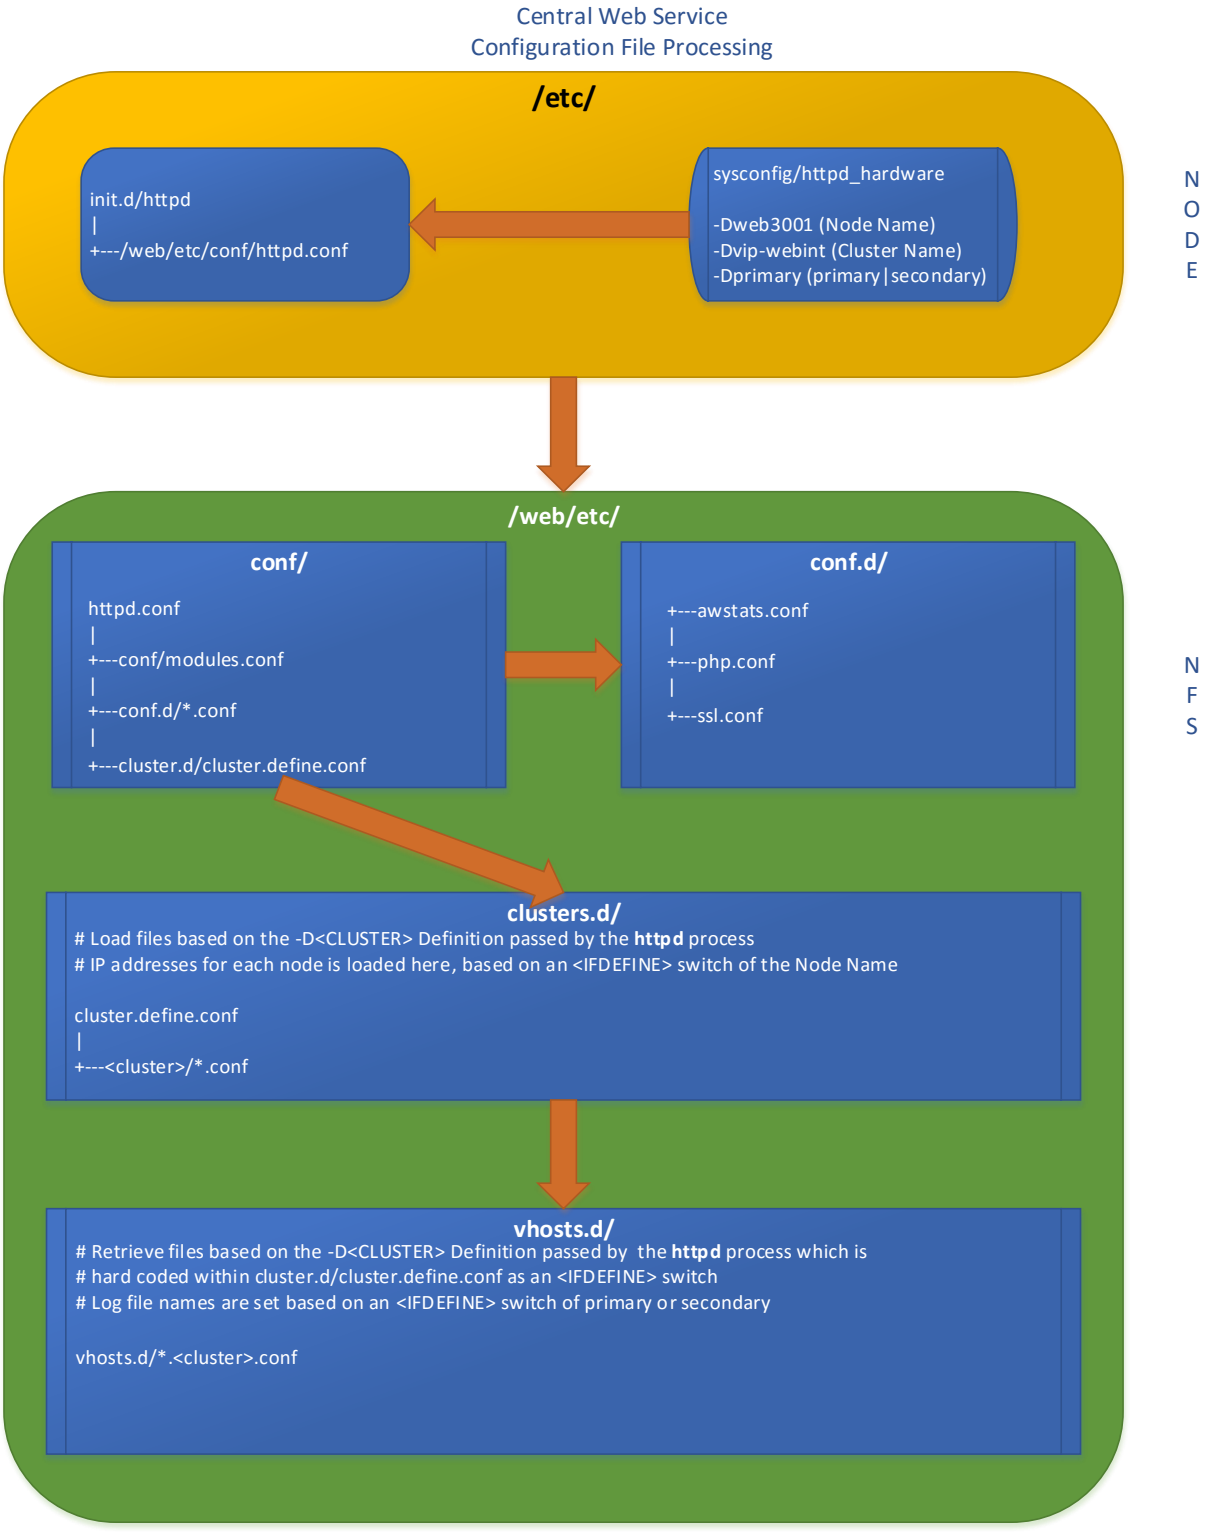

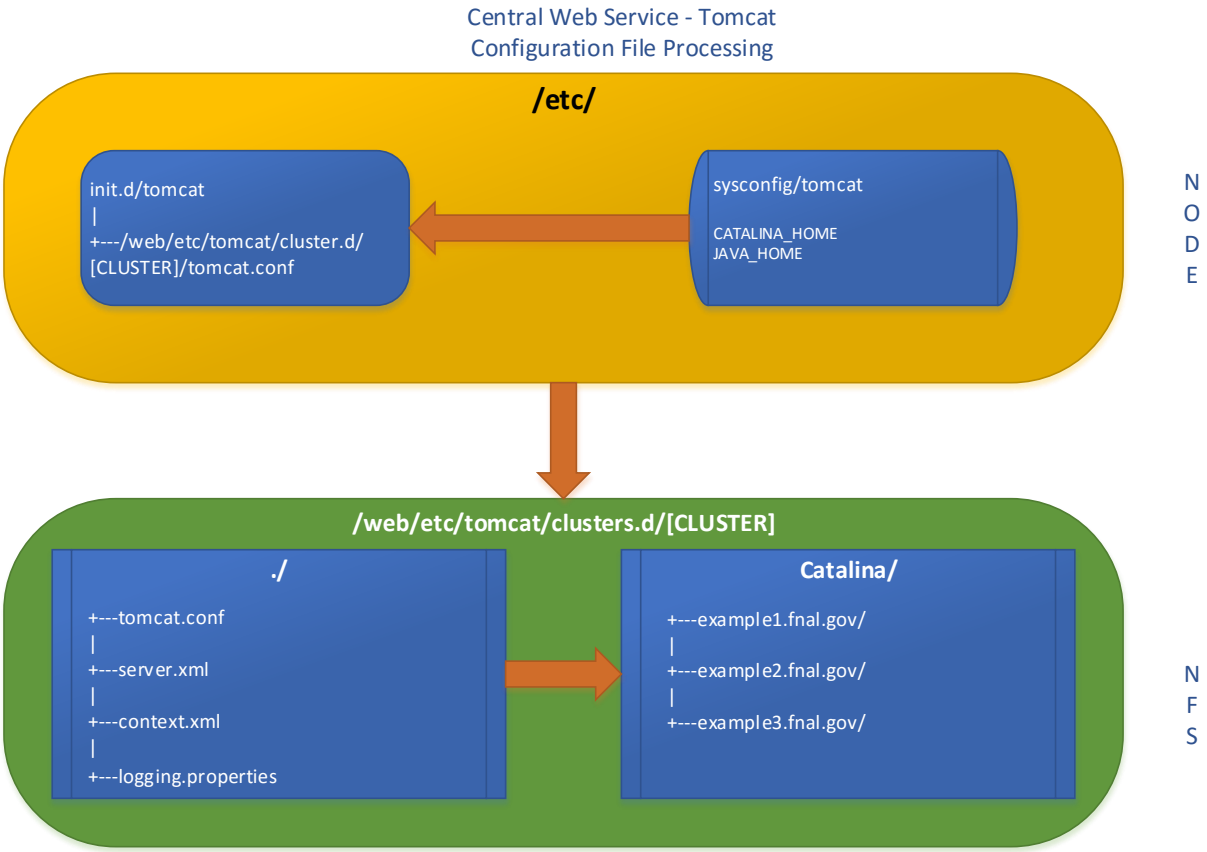
Central Web Service - Tomcat
Configuration File Processing
/etc/
init.d/tomcat
|
+---/web/etc/tomcat/cluster.d/
[CLUSTER]/tomcat.conf
sysconfig/tomcat
CATALINA_HOME
JAVA_HOME
NODE
/web/etc/tomcat/clusters.d/[CLUSTER]
./
+---tomcat.conf
|
+---server.xml
|
+---context.xml
|
+---logging.properties
Catalina/
+---example1.fnal.gov/
|
+---example2.fnal.gov/
|
+---example3.fnal.gov/
NFS

NFSv4 ACL Example

Central Web Service
NFSv4 ACL Overview

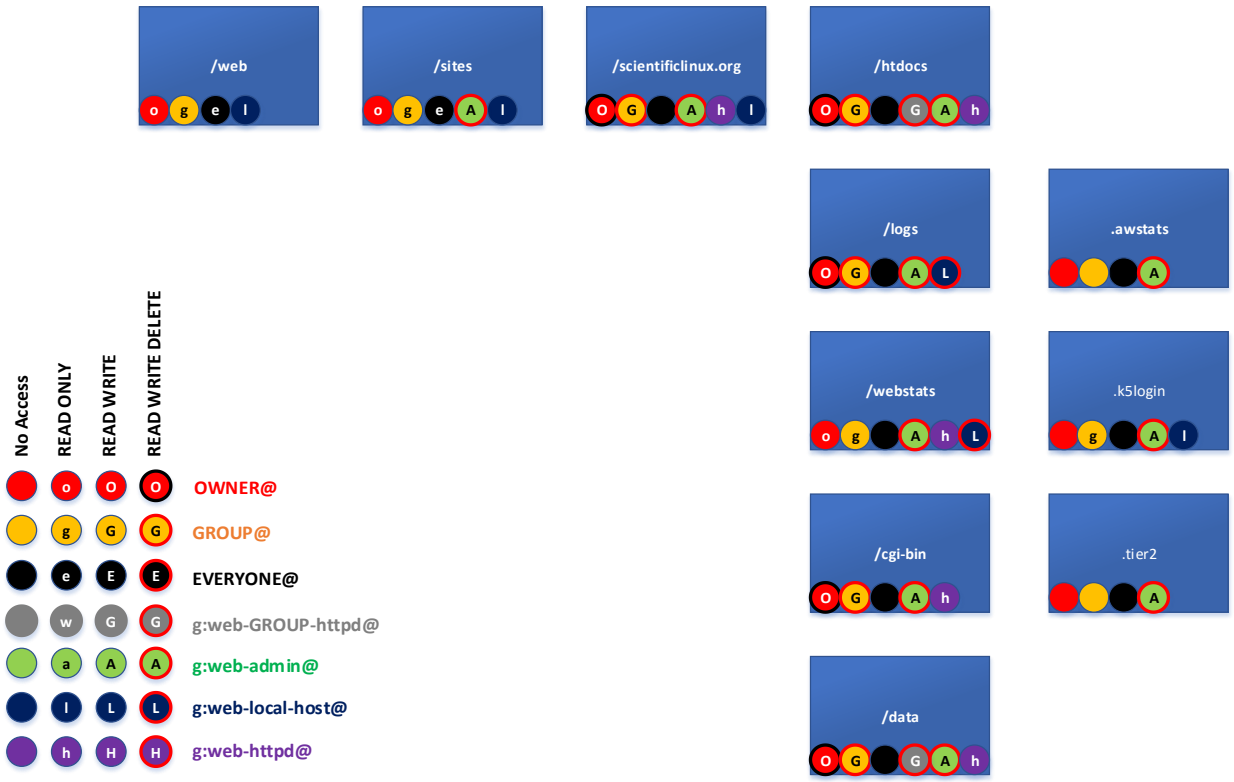

Subnet Allocation

All web servers will be on the 70 subnet. This subnet is reserved exclusively for use by Central Web Hosting.

VLAN ID	Network
VLAN 70 WWW-SUBNET	131.225.70.0/24 2620:6a:0:70::70:0/24
	gw: 131.225.70.1 gw: 2620:6a:0:70::70:1

Web Server Logs

Per standing Computer Security policy, all web server logs are being forwarded to clogger.fnal.gov.

Patching

Patching of the web service is defined in detail in the OLA between the WSA & LSS Groups. In brief, it will be handled in this manner:

- All patches that are of an urgent nature, as identified by the Vendor and/or Computer Security, will be applied immediately.
- All general patches supplied by the vendor will be applied on a monthly basis; Apache httpd, Apache Tomcat, Perl, Python, and PHP are excluded from this general patching cycle and held back a few days to allow the Service Owners to test them on the Tier 3 dev & integration servers.
- When all patches have been cleared for use, the will be applied to the stand-by server first, the server checked to ensure to functions as expected, then applied to the active server.

NFSv4 ACL's

Content access is managed by the use of NFSv4 ACL's. Whereas a POSIX ACL is designed to be one dimensional in that the permissions are fixed to applied to only the owner, the group, and everybody else, NFSv4 ACL's are more three-dimensional. By that we mean that a single file or directory can have multiple owners, group, or hosts that can be granted or denied access, with each of them potentially being different types of access.

The ACL mapping is described like this:

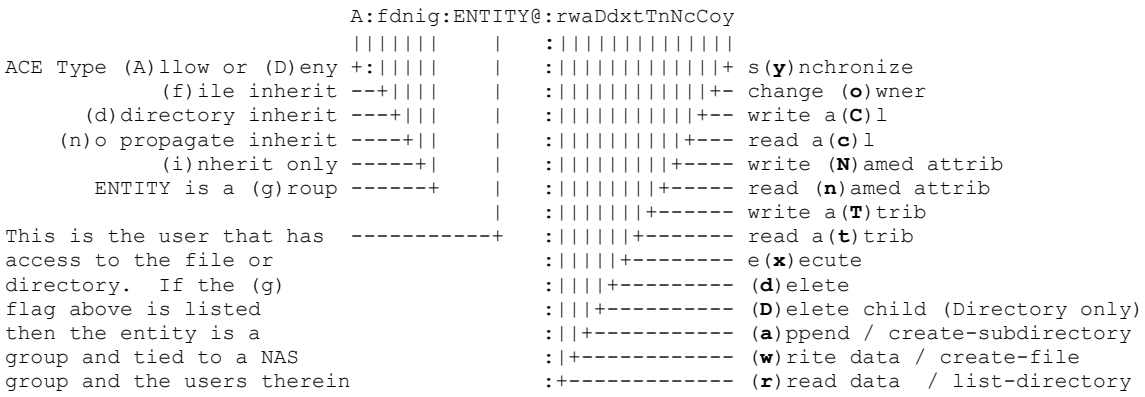

Note: ACL's are "default deny", which means if it is not explicitly granted access via an "A" type ACL, then the entity trying to access the file is denied access. While it is possible to create a "D" type ACL, it is frowned upon as different OS's may interpret the ACL differently. As a result, all ACL's created and used within the NAS are of the "A" type and configured to be least-permissive so that only those entities that need access are granted access.

Example ACL's

The ACL for the root directory of a web vhost looks like this:

```
A:fd:OWNER@:rwaDdxtncy
A:fd:GROUP@:rwaDdxtncy
A:fd:EVERYONE@:tcy
A:fdg:web-admin@fnal.gov:rwaDdxtTnNcCoy
A:fdng:web-httpd@fnal.gov:rxtncy
A:fdng:web-local-host@fnal.gov:rxtncy
A:g:web-tomcat@fnal.gov:rxtncy
```

- The owner of the directory has full rights.
- The group which owns the content, has the same rights.
- Anonymous users (everyone) has no rights to even see the content.

- web-admin, which is a group consisting of members of the ESO/WSA Group, have full rights to the content.
- The apache user, identified by the web-httpd group, has read-only access to the directory.
- The root user, identified by the web-local-host group, has read-only access to the directory.
- The tomcat user, identified by the web-tomcat group, has read-only access to the directory.

The ACL for the logs directory under the directory above looks like this:

```
A:fd:OWNER@:rwaDxtTnNcy
A:fd:GROUP@:rwaDdxtTnNcy
A:fd:EVERYONE@:tcy
A:fdg:web-admin@fnal.gov:rwaDdxtTnNcCoy
A:fdg:web-local-host@fnal.gov:rwadxtTnNcy
A:fdg:web-tomcat@fnal.gov:rwadxtTnNcy
```

- The owner of the directory has full rights.
- The group which owns the content, has the same rights.
- Anonymous users (everyone) has no rights to even see the content.
- web-admin, which is a group consisting of members of the ESO/WSA Group, have full rights to the content.
- The apache user, identified by the web-httpd group, is not listed in this ACL, meaning that it defaults to use the EVERYONE@ ACL, which blocks it access to the directory.
- The root user, identified by the g:web-local-host group, has read/write/delete access to the directory and the files within.
- The tomcat user, identified by the web-tomcat group, has read/write access to the directory.

Information Sensitivity

The data sensitivity on ID is classified in the following table:

	Relative Importance of Protection Needs		
	HIGH (Critical Concern)	**MEDIUM (Important Concern)**	**LOW (Minimum Concern)**
Confidentiality			**X**
Integrity			**X**
Availability			**X**

Risk Identification & Methodology

A threat is the potential for a particular threat-source to successfully exercise a particular vulnerability.

A vulnerability is a weakness that can be accidentally triggered or intentionally exploited.

A threat-source does not present a risk when there is no vulnerability that can be exercised.

Likelihood Determination, Impact Analysis, and Risk Level

The likelihood that each vulnerability will be exploited and the impact of a successful exploit is indicated by the pair of rankings associated with each vulnerability below. Following each vulnerability is the risk level obtained by using the following matrix:

		Impact		
		Low	**Medium**	**High**
Threat Likelihood	**Low**	Low	Low	Low
	Medium	Low	Medium	Medium
	High	Low	Medium	High

Threat Source Identification

There are no threat sources which have not been identified in the Risk Assessment for the General Computing Enclave.

Motivation and Threat Actions

There are no motivations and threat actions which have not been identified in the Risk Assessment for the General Computing Enclave.

Residual Risk Definition

Residual risks are divided into categories based on expected frequency of occurrence after full implementation of all security controls. We consider an occurrence rate to be:

LOW if it is expected to happen <10 times per year
VERY LOW if it is expected to happen <1 time per year
EXTREMELY LOW if it is expected to happen <1 time per five years

Identified Risks, Mitigations, and Residual Risks

General Risks

Apache Tomcat is an Enhanced Offering of the Central Web Hosting Service, built upon Apache httpd as the base service. As such, it is understood that most, if not all, of the Identified Risks, Mitigations, and Residual Risks covered by the Central Web Hosting Risk Assessment found in DocDB Doc.# 5346 will apply here. For the sake of brevity, this document will not repeat those Risks in this document. Instead it will document new Risks, or in the case where a CWH Risk has significantly changed due to the specific configuration differences of the WordPress SaaS offering, identify them here as something new with a reference to the original.

Service Specific Configuration Risks

Risk: The apache process and Tier 7 Servers

Threat & Vulnerability: Tier 7 servers will be granted to ability to have the apache user write and delete content from the htdocs directory. This makes Tier 7 webservers vulnerable to content deletion and defacement.

Mitigation: There are multiple mitigations for this.

- The Tier 7 server is dedicated to websites that have specifically requested Tomcat support. This further lowers the number of possible exposure points.
- Each Tier 7 website will be configured so that only the web servers hosting its content will have the ability to write into the content directory. Should the worst happen and the apache or tomcat user for that server become compromised, the NFSv4 ACL's will limit the damage to only the websites hosted on that server. All other servers that allow read/write to the file system, Tier 2 and future Tier 7 servers, will remain unaffected by the break of a single Tier 7 server.

Residual risk: **VERY LOW**

The Tier 2 servers use a similar configuration. While insecure user code can always make this attack more likely to succeed, we have not seen it used successfully since they were stood up.

Risk: Server Access Points

Threat & Vulnerability: The Tomcat server listens on ports 8080, 8005, and 8009. An attacker connecting to these ports could make requests to applications directly, bypassing HTTP authentication restrictions.

Mitigation: Tomcat will be configured to listen on localhost only. IPTables will be configured to block access to these ports.

Residual risk: **EXTREMELY LOW**
The WordPress SaaS service uses a similar model for PHP-FPM. We have not seen the HTTP authentication bypassed to make direct requests to PHP.

Risk: Proxies and Protected Content

Threat & Vulnerability: Apache httpd is used in a proxy capacity to serve Tomcat content. Proxies can be exploited to bypass normal content protections and access control.

Mitigation: The Tomcat server is configured to listen only on the localhost interface. Apache httpd is configured to proxy only to localhost. If an attacker is able to manipulate requests, they will not be able to see any content not already served by the Tier 7 servers.

Residual risk: **EXTREMELY LOW**
The WordPress SaaS service uses a similar model for PHP-FPM. We have not seen the HTTP proxies used to access unintended content.

Risk: HTTP Header authentication

Threat & Vulnerability: Apache httpd passes authentication user information to Tomcat as an HTTP header. An attacker connecting directly to tomcat could forge authentication data.

Mitigation: The Tomcat server is configured to listen only on the localhost interface. Apache httpd drops incoming authentication header data when proxying to Tomcat, replacing it with its own validated data.

Residual risk: **EXTREMELY LOW**
For this attack to be successful, the Apache configuration would have to be changed unintentionally to allow the forged headers through. We do not make changes to core configuration without testing in development. Additionally, it would be difficult for an attacker to determine which headers to forge.